

SPECIFICATION

Electronic Version 1.2.8

Stylesheet Version 1.0

[Apparatus and method for transparent selection of an Internet server based on geographic location of a user]

Background of Invention

- [0001] The Internet, which is a set of loosely connected computer networks, is growing rapidly in terms of geographical coverage, number of users and number of services. It is a new media that is slowly removing the boundaries between all its users around the world.
- [0002] Applications and services on the Internet are referred by uniform resource locators (URL) such as "ftp://ftp.site/filename", which is an example of a file location in the Internet, or "http://web.site/webpage", which is an example of a web page. However the personal computers (PC) and the servers that host Internet applications communicate by referring to Internet Protocol (IP) addresses of each other, such as 172.168.12.67. PCs hide this complexity by consulting Domain Name Service (DNS), which maps the URLs to IP addresses.
- [0003] Typically when a user wants to access a service or an URL, it first consults its DNS server to resolve the name of the service to an IP address. If the DNS server does not have that information it consults a root DNS server, a gTLD (generic top level domain) server, then finally the DNS server where the service is located. After user's DNS server obtains the IP address of the service it sends this information to the user and then the user accesses the service.
- [0004] Since today's users are located around the world and accessing wide Internet services such as web pages with news or weather forecast, service providers need to

distribute the content at several geographical locations. Today these locations, also known as mirror sites, are updated with the same information.

[0005] In some implementations, while accessing these distributed services, users are asked to choose a mirror site of their choice. In some cases the users may be asked to answer a series of questions such as their country, preferred language, or even the city they live. These steps not only slow the information retrieval process but also frustrate the users. In some other cases, this manual process is automated by the use of "cookies", which are small set of information about the users stored on their PCs and are sent to the servers every time users access the server. This technique obviously has some security and privacy drawbacks.

[0006] Every Internet connected PC and server has a unique IP address assigned by their administrator, or by their Internet Service Provider (ISP). By some methods it is possible to create a database of all of these addresses with their geographical location information. By capturing and examining the user data packets, more specifically the source IP address of DNS queries from these users, and by using the mentioned IP address/location database, it would be possible to find out where the users are located in the world. Therefore by using this technique an Internet service provider could determine the user location and customize its content based on that information.

[0007] This technology can enable ISPs to supply faster services since they can direct the users to the closest servers. Web portal companies can provide the local news, weather forecast or even local advertisements that belong to user's neighborhood. Online shops can restrict certain goods to be sold to certain countries. Similarly, gambling sites would limit their service to some regions where gambling is allowed. Web sites that provide information in different languages would let the user access the information automatically in the language spoken in user's geographical area.

[0008] There are a few implementations of providing content based on users' geo-location. First, it is possible to integrate some software, which has the database of all IP addresses and their locations, to the server. Through this software's application program interface (API) the server determines user's geo-location and modifies or picks the right content based on the user's location. This method obviously has some

limitations as it is highly customized for that service only. It requires customized program development and modification of server configuration. In addition, it supports only the service it is integrated and it needs to be installed on every server. Furthermore, not every type of service can be supported by this method due to limitations in some server programs.

[0009] Another possible method is using specialized network equipment, which is configured to be the authoritative DNS server of the site. All the DNS queries are directed to it. After it determines the location of the user it directs the user to the closest service site relative to user's location. This method not only requires dramatic modification of the DNS and network setup but also limits the capabilities of the DNS since these equipment are not designed to replace DNS servers. For example they do not support Name Server (NS) records, which are used for delegating the sub domains; zone transfers, incremental updates, and other special resource records.

[0010] Therefore a new apparatus and a method are proposed herein to eliminate the shortcomings and limitations of the methods mentioned above while maintaining the DNS functionality with possibly richer feature set.

Summary of Invention

[0011] The invention selects the optimum Internet service relative to the geographical location of the Internet user by transparently modifying DNS messages. In one embodiment, the optimum location is the closest one to the user. Mentioned users and services could be located anywhere in the world.

[0012] In one embodiment, the invention is implemented in a hardware appliance. It also uses a database of IP addresses with their geographical location in the world. This database is created by consulting some public databases (e.g. WHOIS servers) and by some methods that makes latency, routing and distance computations. The database is updated regularly to reflect the modifications in the Internet.

[0013] During the initial setup, every server location that provides the same service around the world is inserted in to the configuration of the invention. The invention then creates a table that maps all the IP addresses from the database to one of the server locations in the world such that the closest server is mapped to each IP

address.

[0014] In the preferred embodiment, the invention is located in front of the existing DNS server that has IP address records (A records) for all the application servers. Every DNS query coming from users passes through the invention. While examining the data packets, the invention determines the IP address of a user's DNS server that sends the query. Then by referring to the previously computed table the invention modifies the DNS messages to direct the user to the closest server.

[0015] The invention works at the second layer of OSI model, examines and modifies only the DNS packets destined to the DNS server of the service. Any other packet is forwarded transparently. This feature of the invention eliminates the need of modifying the configuration of the DNS servers. In addition, it fully supports all DNS operations such as zone transfers, incremental updates, and all types of DNS queries.

[0016] Since the locating the service takes place at the DNS level, the invention can be used for any type of Internet service. It also does not require use of APIs, scripts, custom code development, or any configuration change on the application servers.

Brief Description of Drawings

[0017] Fig. 1 is a network diagram that depicts an example for the standard way of resolving domain names and locating Internet services.

[0018] Fig. 2 is a network diagram that shows an example for the use of the apparatus to select the closest Internet service location for the users.

[0019] Fig. 3 depicts the format of the DNS protocol message.

[0020] Fig. 4 is a flowchart showing the algorithm used by the invention for processing data packets received at the Internet interface.

[0021] Fig. 5 is a flowchart showing the algorithm used by the invention for processing data packets received at the DNS interface.

[0022] Fig. 6 shows an example for a server table that maps user IP addresses to closest server locations.

Detailed Description

[0023] The invention can be implemented in many formats; as a standalone hardware, or as an application integrated with the DNS server on the same platform, or as an application running on the application server. This section provides detailed description of the preferred embodiment, which is the standalone hardware implementation. The detailed functional descriptions herein are representative. The scope of the invention is not limited by this section.

[0024] Fig. 1 shows an example for the standard operation of a domain name resolution. An example user 101 located in New York (NY), which is connected to the Internet through an ISP, wants to access a web site with fully qualified domain name of "www.some-company.com" located in London, UK 104 (the user also could be an employee of a company which is directly connected to the Internet). In order to access this server the user needs to resolve the server's domain name to an IP address. For that, the user's PC sends a DNS query 105 to the user's DNS server 102.

[0025] If the user's DNS server 102 does not have an entry for that domain name in its cache, it sends a query to a root DNS server and a gTLD server to locate the authoritative DNS server. If it does have a cache entry from previous queries it answers the query with that entry. Once it finds the authoritative DNS server the user's DNS server 102 sends a query 106 to the service DNS server 103. The query simply asks the IP address of the service "www.some-company.com". Then the DNS server 103 sends a reply message 107 with the correct IP address of the service. After receiving the reply the user's DNS server 102 sends a reply 108 to the original DNS query (105). Upon determining the IP address of the service the user's PC 101 connects (109) to the web server 104 located in London.

[0026] The previous example illustrates a typical interaction between the PCs, DNS servers and application servers. As the number of users increase or the users get spread around the world, service providers deploy multiple mirror application servers to maintain the service quality. Let's assume that the service in London (Fig.1) is duplicated in Washington, D.C. and the invention will be used to select the closest server location.

[0027] Fig. 2 illustrates this expansion of service coverage. Besides the original web

server 206 located in London there is a new mirror server located in DC 205. In addition, the invention 203 is located in front of the DNS server 204 in London.

[0028] In order to resolve the IP address of the web server 206, the user in NY 201 consults its DNS server 202 also located in NY. The user 201 sends a DNS query 209 to its DNS server 202. After determining the authoritative DNS server, the DNS server 202 sends a query 210 towards the DNS server in London 204 (assuming there is no cache entry for that domain in the DNS server 202).

[0029] The invention 203 is located in front of the DNS server 204 such that every data packet sent to the DNS server 204 goes through it. The invention determines the packet type by parsing the proper fields in the data packets such as the destination IP address, the destination port number, or the DNS message type, etc. If the invention captures a DNS query packet, which is sent to the DNS server 204, it examines the content of the DNS message for further processing. If the packet is not a DNS packet, it is forwarded to the DNS server.

[0030] Fig. 3 depicts the format of a DNS message. If the QR field 302 in the DNS message is 0 (zero) it indicates that the message is a DNS query. The question section 314 actually has the queried domain name. The answer section 315 is left blank in the query but populated with a proper answer in the response message in which QR 302 is set to 1 (one). The invention 203 also processes other fields to provide correct operation of the DNS service. For example, since each DNS response should have the same DNS Query Number with the query the invention keeps track of each DNS query-response pair by recording the DNS Query Number 301. Also every response message should carry the original Question Section 314 besides the Answer Section 315.

[0031] In the preferred embodiment, the invention 203 has two network interfaces: one of them is connected to the DNS server 204 (DNS interface), while the other is connected to the backbone network through which the Internet is accessed (Internet interface).

[0032] Fig. 4 shows the algorithm utilized when a packet is received at the Internet interface of the invention. When a packet is received 401, the invention examines whether it is sent to the DNS server 402, or whether it is a DNS query 404. If the

answer is no to any of these questions the packet is forwarded to the DNS interface 403 and sent to the DNS server 204. If it is a DNS query, the Question Section 314 of the DNS message is examined 405. Then this data is checked whether it is the domain name under question 406 (localized query or not). If it is not a localized query then the message is again forwarded to the DNS server interface 403.

[0033] If it is a localized query, which is the DNS query for the domain name under question, then the source IP address of the DNS query is determined (407). It is assumed that DNS server 202 is in close proximity of the user 201. After that the invention looks up the IP address/server table to choose the closest server location according to the source IP address of the DNS query. Then the invention modifies the DNS query's question section with the server location. Referring to the example above, the original DNS query, which was asking the IP address of "www.some-company.com", is converted into "dc.some-service.com" 408. After that the packet is forwarded (403) to the DNS server 204.

[0034] Similarly, fig. 5 depicts the algorithm utilized when a packet is received at the DNS interface of the invention. When a packet is received 501, the invention examines whether it is a DNS response 502. If the packet is not a DNS response, it is forwarded to the Internet interface 503. If it is a DNS response, the invention parses it 504 to check whether the message is a response to a previously processed localized message 505. If the packet is not a localized response message it is again forwarded to the Internet interface 503. If it is a response to a localized query then the Question Section 314 of the response is changed to the original one, for the example mentioned above, from "dc.some-service.com" to "www.some-company.com" 506. Then the response is forwarded to the Internet interface 503.

[0035] Referring back to fig. 2, the invention 203 modifies the DNS query 210 according to the geo-location of the user's DNS server 202. Then it forwards the modified query to the web server's DNS server 204, which has the address records for all the web servers (211). When the DNS server 204 sends a response 212, the invention captures it again. After that the invention modifies the response message accordingly and sends it back (213) to the user's DNS server 202. The user getting the response 214 from its DNS server 202 easily connects (215) to the closer server site in DC 205.

Similarly the user located in Frankfurt 207 chooses to connect (208) the server in London 206 due to its closer location.

[0036] Fig. 6 shows a portion of the lookup table used by the invention. The first column 601 has all the possible IP addresses on the Internet. The second column 602 has the corresponding geo-locations which were obtained from the IP address database, while the third one 603 has the assigned servers for each IP address. For example, for the IP address 610, which belongs to the example user DNS server 202, the invention knows that (from the database) the address is located in New York 611. In addition, the invention has assigned the address to the server in DC 612. Therefore the invention modifies the DNS query coming from the user DNS server 202 for the web service 206 in such a way that the DNS response contains the IP address of the server located in DC 612. Hence the user in NY accesses the server in DC 205, which is closer than the one in London 206. Similarly the user 207 whose DNS server's IP address is "10.175.183.32"604 which is located in Frankfurt 605 accesses the server in London (606), which is closer than the one in DC 204.

[0037] The lookup table is created from two inputs; the database that lists all Internet addresses with their geo-locations and the list of geo-locations of the application servers. The IP database is obtained externally and the invention has an appropriate interface to it. Both the IP database and the list of geo-locations of the servers are entered into the invention during the initial configuration. The size and the content of the lookup table are dependent on the service and the distribution of the users. For example, if the invention is going to be used for a web server for Europe only then there is no need to include IP addresses that belongs to Asia.

[0038] In one embodiment, the lookup table can be formed by dividing the IP database into equally sized regions. The number of regions would be the same as the number of servers such that each server would be assigned to a region. Any IP address in a region would then be assigned to the server in that region.

[0039] In a further embodiment, servers can be assigned weights so that the ones with higher weights would cover larger geographical area. This feature can be used for areas where the Internet user distribution is less populated.

[illegible]

[0041] Another embodiment of the invention can check the health of the application servers by some methods so that if a server becomes unavailable then the user IP addresses assigned to this server would be re-assigned to the next closest server. When this feature is utilized the invention can provide a very reliable service where the users can always access the service as long as there is at least one server available.